

Network Working Group
Internet-Draft
Updates: [3315](#),3633 (if approved)
Intended status: Standards Track
Expires: July 5, 2014

O. Troan
B. Volz
Cisco Systems, Inc.
January 1, 2014

Issues with multiple stateful DHCPv6 options
draft-ietf-dhc-dhcpv6-stateful-issues-05.txt

Abstract

Dynamic Host Configuration Protocol for IPv6 (DHCPv6) was not written with the expectation that additional stateful DHCPv6 options would be developed. IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6 shoe-horned the new options for Prefix Delegation into DHCPv6. Implementation experience of the CPE model described in [RFC6204](#) has shown multiple issues with the DHCPv6 protocol in supporting multiple stateful options. This document updates [RFC3315](#) and indirectly [RFC3633](#).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 5, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

| | | |
|----------------------|-----------------------------------------------------------------|--------------------|
| 1. | Introduction | 2 |
| 2. | Conventions | 3 |
| 3. | Terminology | 3 |
| 4. | Handling of multiple IA options types | 3 |
| 4.1. | Advertise Message | 4 |
| 4.2. | Placement of Status Codes | 5 |
| 4.3. | T1/T2 Timers | 5 |
| 4.4. | Renew Message | 6 |
| 4.5. | Rebind Message | 7 |
| 4.6. | Confirm Message | 8 |
| 4.7. | Release Message | 9 |
| 4.8. | Multiple Provisioning Domains | 9 |
| 5. | IANA Considerations | 10 |
| 6. | Security Considerations | 10 |
| 7. | Acknowledgements | 10 |
| 8. | References | 10 |
| 8.1. | Normative References | 10 |
| 8.2. | Informative References | 10 |
| | Authors' Addresses | 10 |

[1.](#) Introduction

DHCPv6 [[RFC3315](#)] was not written with the expectation that additional stateful DHCPv6 options would be developed. DHCPv6 Prefix Delegation [[RFC3633](#)] shoe-horned the new options for Prefix Delegation into DHCPv6. Implementation experience of the CPE model described in [[RFC6204](#)] has shown multiple issues with the DHCPv6 protocol in supporting multiple stateful options.

This document describes a number of problems encountered with multiple IA option types into DHCP and recommended changes to the DHCPv6 protocol specifications.

The intention of this work is to modify the DHCP protocol specification to support multiple IA option types within a single DHCP session. This problem can also be solved by implementing a separate DHCP session (separate client state machine) per IA option type. This latter approach has a number of issues: additional DHCP protocol traffic, 'collisions' between stateless options also included with the IA options, divergence in that each IA option type specification specifies its 'own' version of the DHCP protocol.

The changes described in this document will be incorporated in a new revision of the DHCPv6 protocol specification [[RFC3315](#)].

2. Conventions

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

3. Terminology

Stateful options Options that require dynamic binding state per client on the server.

Identity association (IA): A collection of stateful options assigned to a client. Each IA has an associated IAID. A client may have more than one IA assigned to it; for example, one for each of its interfaces. Each IA holds one type of IA option; for example, an identity association for temporary addresses (IA_TA) holds temporary addresses (see "identity association for temporary addresses"). Throughout this document, "IA" is used to refer to an identity association without identifying the type of stateful option in the IA.

4. Handling of multiple IA options types

DHCPv6 was written with the assumption that the only stateful options were for assigning addresses. DHCPv6 PD describes how to extend the DHCPv6 protocol to handle prefix delegation, but [[RFC3633](#)] did not consider how DHCP address assignment and prefix delegation could co-exist.

If a client requests multiple IA option types, but the server is configured to only offer a subset of them, the client could react in several ways. Reset the state machine and continue to send Solicit messages, create separate DHCP sessions for each IA option type and continue to Solicit for the unfulfilled IA options, or it could continue with the single session, and include the unfulfilled IA options on subsequent messages to the server.

Proposed solution: the client should keep a single session with the server and include the missing options on subsequent messages (Request, Renew, and Rebind) to the server.

4.1. Advertise Message

[RFC3315] specifies that a client must ignore an Advertise message if a server will not assign any addresses to a client. A client requesting both IA_NA and IA_PD, with a server that only offers one of them, is not supported in the current protocol specification.

Proposed solution: a client SHOULD accept Advertise messages, even when not all IA option types are being offered. A client SHOULD ignore an Advertise message when no bindings at all are being offered. The client SHOULD include the not offered IA option types in its Request.

Replace [Section 17.1.3 of \[RFC3315\]](#): (existing errata)

The client MUST ignore any Advertise message that includes a Status Code option containing the value NoAddrsAvail, with the exception that the client MAY display the associated status message(s) to the user.

With:

The client MUST ignore any Advertise message that contains no bindings (if only IA_NA and/or IA_TA options were requested, this is a message that includes a Status Code option containing the value NoAddrsAvail), with the exception that the client MAY display the associated status message(s) to the user.

And, replace:

- The client MAY choose a less-preferred server if that server has a better set of advertised parameters, such as the available addresses advertised in IAs.

With:

- The client MAY choose a less-preferred server if that server has a better set of advertised parameters, such as the available options advertised in IAs.

It is important to note that the receipt of a Advertisement without any bindings does not imply that the client should restart the Solicit retransmissions timers. Doing so would lead to a Solicit/Advertisement storm.

4.2. Placement of Status Codes

In Reply messages IA specific status codes (i.e., NoAddrsAvail, NotOnlink, NoBinding, NoPrefixAvail) are encapsulated in the IA option. In Advertisement messages the Status Code option with the NoAddrsAvail code is in the "global" scope. That makes sense when the failure case is fatal. With the introduction of multiple IA option types, there might be a case where a server is not willing to offer addresses, but might be willing to offer other stateful option types.

While a Status Code option is implicitly bound to a specific type of IA, e.g. NoPrefixAvail is only applicable to IA_PD and NoAddrsAvail is only applicable to IA_NA/IA_TA, it may be problematic to make this assumption for all status codes. Ideally the Status Code option should be encapsulated in the IA option for all DHCP messages. This makes Advertisement messages equal to Reply messages.

Proposed solution: No change. For backwards compatibility, the NoAddrsAvail Status Code option when no addresses are available will be kept in the global scope for Advertise messages. Other IA option types MUST encapsulate the Status Code option within the IA option.

To clarify further: when a client requests both IA_NA and IA_PD, and the server can offer IA_PD but not IA_NA, the server sends an Advertise response containing no IA_NA option, a status code option of NoAddrsAvail, and one or more IA_PD options containing IAPREFIX options.

4.3. T1/T2 Timers

The T1 and T2 timers determine when the client will contact the server to extend lifetimes of information received in an IA. How should a client handle the case where multiple IA options have different T1 and T2 timers?

In a multiple IA option types model, the T1/T2 timers are protocol timers, that should be independent of the IA options themselves. If we were to redo the DHCP protocol from scratch the T1/T2 timers should be carried in a separate DHCP option.

Proposed solution: The server SHOULD set the T1/T2 timers in all IA options in Reply and Advertise messages to the same value. To deal with the case where servers have not yet been updated to do that, clients MUST use the shortest (explicit or implicit) T1/T2 timer (larger than 0) in any IA options in the Reply. Longer T1/T2 timers are ignored.

4.4. Renew Message

The Renew message, as described in [[RFC3315](#)], allows a client to only renew bindings assigned via a Request message.

In a multiple IA option type model, the Renew does not support the ability for the client to renew one IA option type while requesting bindings for other IA option types that were not available when the client sent the Request.

Proposed solution: The client should continue with the IA options received, while continuing to include the other IA options in subsequent messages to the server. The client and server processing need to be modified. Note that this change makes the server's IA processing of Renew and Rebind similar to the Request processing.

Replace [Section 18.1.3 of \[RFC3315\]](#):

At time T1 for an IA, the client initiates a Renew/Reply message exchange to extend the lifetimes on any addresses in the IA. The client includes an IA option with all addresses currently assigned to the IA in its Renew message.

With:

At time T1 for an IA, the client initiates a Renew/Reply message exchange to extend the lifetimes on any addresses in the IA. The client includes an IA option with all addresses currently assigned to the IA in its Renew message. The client also includes an IA option for each binding it desires but has been unable to obtain.

Replace [Section 18.2.3 of \[RFC3315\]](#):

If the server cannot find a client entry for the IA the server returns the IA containing no addresses with a Status Code option set to NoBinding in the Reply message.

With:

If the server cannot find a client entry for the IA the server creates the bindings for that client according to the server's policy and configuration information and records the IAs and other information requested by the client.

Note that clients that communicate with servers that do not support this updated Renew processing will receive the NoBinding status for the IA which had no bindings. The client MUST continue to process the other IAs in the Reply. The client MAY attempt a Solicit/

Advertise/Request/Reply sequence periodically to obtain bindings for these IAs. However, it MUST limit the frequency at which it does this to no more often than the renewal frequency.

4.5. Rebind Message

The Rebind message, as described in [\[RFC3315\]](#) does not explicitly specify what a server should do when an IA option which contains no addresses is present.

Replace [Section 18.1.4 of \[RFC3315\]](#):

At time T2 for an IA (which will only be reached if the server to which the Renew message was sent at time T1 has not responded), the client initiates a Rebind/Reply message exchange with any available server. The client includes an IA option with all addresses currently assigned to the IA in its Rebind message.

With:

At time T2 for an IA (which will only be reached if the server to which the Renew message was sent at time T1 has not responded), the client initiates a Rebind/Reply message exchange with any available server. The client includes an IA option with all addresses currently assigned to the IA in its Rebind message. The client also includes an IA option for each binding it desires but has been unable to obtain.

Replace [Section 18.2.4 of \[RFC3315\]](#) with the following text to clarify how the server should handle all of the possible conditions:

When the server receives a Rebind message that contains an IA option from a client, it locates the client's binding and verifies that the information in the IA from the client matches the information stored for that client.

If the server finds the addresses in the IA for the client and the server determines that the addresses in the IA are appropriate for the link to which the client's interface is attached according to the server's explicit configuration information, the server SHOULD send back the IA to the client with new lifetimes and T1/T2 times.

If the server cannot find a client entry for the IA and the server determines that the addresses in the IA are not appropriate for the link to which the client's interface is attached according to the server's explicit configuration information, the server MAY send a Reply message to the client containing the client's IA, with the lifetimes for the addresses in the IA set to zero. This Reply

constitutes an explicit notification to the client that the addresses in the IA are no longer valid. In this situation, if the server does not send a Reply message it silently discards the Rebind message.

If the server cannot find a client entry for the IA and the server determines that the addresses in the IA are appropriate for the link to which the client's interface is attached according to the server's explicit configuration information, and the addresses are not in use, the server MAY assign the addresses to the client and send a Reply message to the client with new lifetimes and T1/T2 times for the bindings. If the server cannot assign the addresses to the client, the server returns the IA containing no addresses with a Status Code option set to NoBinding in the Reply message.

Note: A server SHOULD NOT provide additional bindings to the client as the client could accept a Reply from a different server (this is the same issue as in the Discussion under the Rapid Commit option, see [section 22.14](#)).

The server constructs a Reply message by setting the "msg-type" field to REPLY, and copying the transaction ID from the Rebind message into the transaction-id field.

The server MUST include a Server Identifier option containing the server's DUID and the Client Identifier option from the Rebind message in the Reply message.

The server includes other options containing configuration information to be returned to the client as described in [section 18.2](#).

4.6. Confirm Message

The Confirm message, as described in [[RFC3315](#)], is specific to address assignment. It allows a server without a binding reply to the message, under the assumption that the server only needs knowledge about the prefix(es) on the link, to inform the client that the address is likely valid or not. This message is sent when e.g. the client has moved and needs to validate its addresses. Not all bindings can be validated by servers and the Confirm message provides for this by specifying that a server that is unable to determine the on-link status MUST NOT send a Reply.

Note: Confirm has a specific meaning and does not overload Renew/Rebind. It also is lower processing cost as the server does NOT need to extend lease times or otherwise send back other configuration options.

Proposed solution: Allow and specify the Confirm message for other IA option types. The server performs the same test as for addresses on the delegated prefixes (see [\[RFC3315\]](#), [section 18.2.2](#)).

Replace [Section 12.1 of \[RFC3633\]](#):

If such verification is needed the requesting router MUST initiate a Rebind/Reply message exchange as described in [section 18.1.4](#), "Creation and Transmission of Rebind Messages" of [RFC 3315](#), with the exception that the retransmission parameters should be set as for the Confirm message, described in [section 18.1.2](#), "Creation and Transmission of Confirm Messages" of [RFC 3315](#). The requesting router includes any IA_PDs, along with prefixes associated with those IA_PDs in its Rebind message.

...

The Confirm and Decline message types are not used with Prefix Delegation.

With:

If such verification is needed the requesting router MUST initiate a Confirm message exchange as described in [section 18.1.2](#), "Creation and Transmission of Confirm Messages" of [RFC 3315](#). The requesting router includes any IA_PDs, along with prefixes associated with those IA_PDs in its Confirm message.

...

The Decline message type is not used with Prefix Delegation.

[4.7.](#) Release Message

A client can release any individual lease at any time. A client can get "back" a lease by using a Renew message. It MAY do this at any time, though must avoid creating a Renew storm. E.g. wait until T1.

[4.8.](#) Multiple Provisioning Domains

This document has assumed that all DHCP servers on a network are in a single provisioning domain and thus should be "equal" in the service that they offer. This was also assumed by [\[RFC3315\]](#) and [\[RFC3633\]](#).

One could envision a network where the DHCP servers are in multiple provisioning domains, and it may be desirable to have the DHCP client obtain different IA types from different provisioning domains. How a client detects the multiple provisioning domains and how it would

interact with the multiple servers in these different domains is outside the scope of this document.

5. IANA Considerations

This specification does not require any IANA actions.

6. Security Considerations

There are no new security considerations pertaining to this document.

7. Acknowledgements

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.
- [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", [RFC 3633](#), December 2003.

8.2. Informative References

- [RFC6204] Singh, H., Beebee, W., Donley, C., Stark, B., and O. Troan, "Basic Requirements for IPv6 Customer Edge Routers", [RFC 6204](#), April 2011.

Authors' Addresses

Ole Troan
Cisco Systems, Inc.
Philip Pedersens vei 20
N-1324 Lysaker
Norway

Email: ot@cisco.com

Bernie Volz
Cisco Systems, Inc.
1414 Massachusetts Ave
Boxborough, MA 01719
USA

Email: volz@cisco.com